

CS 268 1st Term Exam
Date: March 16, 2005

Name:
SID:

Problem	Points
1	/ 16
2	/ 12
3	/ 20
4	/ 20
5	/ 12
6	/ 20
Total	/ 100

Please be concise.

1) End-to-end Argument (16 pts)

(a) State the end-to-end argument in three (or fewer) sentences.

The quote from the paper is: "The function in question can completely and correctly be implemented only with the knowledge and help of the application standing at the end points of the communication system. Therefore, providing that questioned function as a feature of the communication system itself is not possible. (Sometimes an incomplete version of the function provided by the communication system may be useful as a performance enhancement.)" This can be interpreted in many ways, but all of them urge caution when placing functionality at lower layers.

(b) Using the E2E argument, take a position on whether multicast belongs in the network layer and defend it in three or fewer sentences.

Almost anything was acceptable here: in-network multicast is far more efficient, but it can be done by ends.

(c) Using the E2E argument, take a position on whether mobility (as in mobile IP) belongs in the network layer and defend it in three or fewer sentences.

Almost anything was acceptable here: in-network mobility has more efficient routing, but application-layer mobility is able to adapt to application-specific needs.

(d) Using the E2E argument, take a position on whether Quality of Service belongs in the network layer and defend it in three or fewer sentences.

Almost anything was acceptable here, but many people argued for putting QoS at the application-layer without explaining how that could be done.

2) TCP (12 pts)

TCP uses AIMD to adjust window sizes. Why aren't the following choices used? (Describe a problem with each one, no more than a two sentences for each.)

a) MIMD

Stable cycle is unfair.

b) MIAD

In simple fixed loss probability, this has unstable fixed point. It can also result in bandwidth capture, where one flow gets all the bandwidth.

c) AIAD

Stable cycle is unfair.

3) Bandwidth Estimation (20 pts)

In the following, ignore all issues related to cross-traffic or other sources of noise.

(a) How can packet pairs determine the bandwidth of a bottleneck link?

Send two packets back-to-back. The delay between their arrival at the receiver is P/C where P is packet size and C is bottleneck bandwidth.

(b) How can variable sized packet probes measure the bandwidth of every link along the path?

Delay at a router i is $D_i = A_i + P/C_i$. The delay for the first k routers is:

$$S(k) = \sum_{i=1}^k D_i$$

Measure delays to first k routers with varying P and varying k (using TTL). Fit to values of A_i and C_i .

4) Congestion control (20 pts)

(a) Name two problems of TCP that XCP solves, and briefly explain (use no more than three sentences for each case) how it does solve these problems.

Fast convergence: XCP can find the available bandwidth via explicit signaling in one RTT. In contrast, to find the available bandwidth, TCP needs to probe the network using the slow start and additive increase, which will take multiple RTTs (especially in high bandwidth-delay product networks).

Fairness: XCP achieves max-min fairness by having the end-hosts report the estimated RTTs. In contrast, TCP implements proportional fairness where long flows can receive significantly less bandwidth.

(b) Does TFRC have an advantage over TCP for small flows, i.e., flows consisting of a couple of packets? Use no more than three sentences to explain your answer.

No. TFRC still uses slow-start initially. If the flow has very few packets, TFRC won't be able to estimate the loss rate, p .

5) Router Architecture (12 pts)

Consider a router with N inputs and N outputs. The router has no buffers (either at the inputs or outputs), and has a speedup of **one**, i.e., an input can send data to at most one output and an output can receive data from at most one input. Assume that each input has always packets to send, and that each packet is forwarded to an output with equal probability. Compute the expected utilization of the router. (The expected utilization is the ratio between the expected number of packets that are forwarded and the router capacity.)

We first compute the expected number of outputs $E(\text{empty_outputs})$ that do not receive any packets.

Consider an output i . The probability that a given input j does **not** send a packet to output i is $(N-1)/N$. Thus the probability that none of the inputs send a packet to output i is $(1 - 1/N)^N$. By linearity of expectation, we have


$$E(\text{empty_outputs}) = N \cdot (1 - 1/N)^N$$

Finally, the utilization is the number of outputs that have packets to send divided by the total number of outputs:

$$U = (N - E(\text{empty_outputs})) / N = 1 - (1 - 1/N)^N$$

6) Routing Protocols (20 pts)

Consider the following network topology in which the black node claims to have a link to A of cost one. Consider two routing protocols, link-state and distance-vector, that compute shortest path on this network. Assume that all links in the network are bidirectional and have a cost of one. Answer the following questions:



(a) Is it possible for routers using the distance-vector protocol to filter out the bogus information provided by D? Explain your answer using no more than three sentences.

No. With distance-vector, each node maintains for each destination only the next hop and the distance to that hop. Since a node doesn't know the path to the destination, in general it cannot discover the discrepancy between the info provided by D and A. For example, node C has no choice but to believe D that it has a path to A.

(b) Is it possible for routers using the link-state protocol to filter out the bogus information provided by A? Explain your answer using no more than three sentences.

Yes. With link-state, every node eventually knows the entire topology. This way a node will see that D claims that it has a link to A, while A doesn't, so it can remove the bogus link.

Note: Consider a second scenario in which there is indeed a link between A and D, but A claims incorrectly that it doesn't have a link to D. There is no way for a node to differentiate between this scenario and the previous one (in which D misbehaves). However, the only thing that A will accomplish in the 2nd scenario is to decrease the number of packets routed through it, which is contrary to the typical goal of an attacker to attract more traffic.